Web-enabled Automated Acquisition: Design and Experiments

EARP Project Final Report

PI:

Olivia Sheng, Professor

Department of Management Information Systems

University of Arizona, Tucson, Arizona 85721

co-PI:

Daniel Zeng, Assistant Professor

Department of Management Information Systems

University of Arizona, Tucson, Arizona 85721

Research Category: Exploratory

Abstract

Advances in information technologies and the electronic marketplace present significant opportunities for procurement management, calling for new paradigms for developing and evaluating Web-enabled intelligent procurement decision support systems. This paper examines issues relating to the design and evaluation of advanced knowledge-driven software infrastructures based on agent technology for online procurement. We discuss various system architecture decisions relevant to developing such an agent-based procurement system and report a case study concerning the implementation of a smart trading strategy in the context of a publicly-held agent trading competition. In this report, we also evaluate the effectiveness and efficiency of new online acquisition technologies using experimental economics, and advocate the importance of studying the interactions among intelligent procurement systems and their human users in digital marketplaces. We present an experimental study designed for studying these interactions for a simple online procurement mechanism based on English auctions and discuss resulting experimental findings.

1 Introduction

Procurement management is concerned with coordinating all activities pertaining to purchasing goods and services necessary to accomplish the mission of an enterprise. Recent years have seen increasing awareness of the importance of procurement in both private and public sectors, partially driven by recent trends on lean manufacturing and supply chain integration, and by the emergence of the new digital economy enabled by advances in networking and Internet-related technologies.

In recent years, two major types of electronic commerce, business-to-consumer (B2C) and business-to-business (B2B) e-commerce, have both been experiencing tremendous growth. Despite the recent slow-down in its growth, e-commerce continues to spring up to industry-specific markets and has been accepted as a significant part of the consumer market [33, 22].

Procurement management is one of the B2B areas in which the impact of e-commerce is significant and yet the specific opportunities and challenges are to be identified and explored. One of the main lessons learned by e-commerce participants and advocates is that although IT is a critical enabling factor, the economic and business aspects of e-commerce and related industrial procurement behaviors have to be seriously studied in the context of IT-enabled online economic institutions. Some specific questions of both scientific and practical significance that need to be answered are:

- Many forms of online economic institutions have already been implemented and adopted including different variations of the posted-price market and one-sided and two-sided auctions (e.g., eBay, Yahoo Auctions, and Fastparts.com). How efficient and effective are these new institutions in practice?

- Recent years have witnessed the emergence of intelligent software (e.g., ADEPT [7], eMediator [25], and Kasbah [2]) that offers sophisticated decision support and automates certain aspects of online economic transactions. More in-depth research is called for to guide the development of such automated trading mechanisms in new and emerging online institutions. Will the adoption of these automated mechanisms change their user's economic and social behavior? What are the impacts of these technologies from the economic, societal, and organizational perspectives?

Our research is motivated to address these fundamental challenges and research issues by designing and implementing a flexible online procurement experimental infrastructure and associated automated trading mechanisms, and by evaluating these procurement models and technologies through experimental research methodologies. In general, our work is aimed at delivering solution concepts and sound scientific evaluation to help procurement managers achieve more efficient and effective procurement management through the use of IT and e-commerce. Our specific approach is (a) identifying problem areas in procurement in which such advanced IT as software agents can be deployed to provide value-added information and decision support, (b) developing advanced IT (including both enabling infrastructures and domain-specific computational decision support) tailored for procurement-related challenges, and (c) evaluating these technologies in experimentally controlled procurement environments guided by the science of experimental economics.

In this paper, we report our research efforts on developing and evaluating advanced procurement infrastructures and related decision making frameworks based on software agents. The rest of the paper is organized as follows. In Section 2, we present a brief overview of software agent technology and point out several major ways through which software agents can facilitate procurement management. Section 3 presents general experimental design and underlying methodology for evaluating advanced IT in strategic interactions such as procurement. In Section 4, we focus on auction-based procurement management and discuss related research issues. Section 5 presents the system architecture of an agent-enabled procurement testbed developed as part of our EARP effort. In Section 6, we report an experimental study that examines human-agent interaction in Internet-based English auctions using this testbed. Section 7 presents the development of an intelligent trading strategy, implemented as a software agent, in the context of Trading Agent Competition (TAC), an agent-based auction tournament organized by the Multi-Agent Systems research community. We conclude the paper in Section 8 by summarizing our research findings and presenting future research.

2 Software Agent-Facilitated Procurement Management

Software agent technology is one of the fastest growing areas of research and system development in Information Technology [8]. Software agents represent both a new modeling framework and a new system development paradigm for many complex information management and decision tasks [23, 9]. Large-scale manufacturing scheduling and supply chain management are example application areas where active agent-based system development efforts are in place and proved to be highly successful [21, 30, 11, 34, 20, 26].

Artificial Intelligence, Software Engineering, and Human-Computer Interface are among the main academic disciplines that have contributed to the development of software agent technology [8]. In a relatively new and rapidly growing multi-disciplinary field such as software agent, it is not surprising that many agent-related concepts and frameworks are confusing, misused, and incoherent. Even the definition of an agent is still being debated [27, 36]. In our research, we view agents in a relatively broad sense as defined in [8]. For us, an agent is a computer program that is situated in some environment which may include other agents, and is capable of autonomous action in order to meet its user-delegated objectives. We list below three major categories of agent-based computational systems. This categorization by no means represents a systematic taxonomy of all existing agent applications. Nevertheless, it offers a conceptual reference framework for the reported research and a majority of agent applications indeed can be characterized in this categorization scheme.

1. Most well-known agent applications are in the area of *information management*, often associated with the Internet. According to the current working notion, information management agents are programs that act on behalf of their human users in order to perform laborious information gathering tasks, such as locating and accessing information from various on-line information sources, resolving inconsistencies in the retrieved information, filtering away irrelevant or unwanted information, integrating information from heterogeneous information sources, and adapting over time to their human users' information needs and the shape of the Infosphere. For instance, many agent-oriented approaches have focused on *interface agents*—a single agent with simple knowledge and problem solving capabilities whose main task is information filtering to alleviate the user's cognitive overload [15, 19]. Another type of agent is the *SoftBot* [4], a single agent with general knowledge that performs a wide range of user-delegated information-finding tasks.

   Recent years have seen the rise and quick adoption of multi-agent based approaches for information management tasks, addressing several limitations of centralized single-agent approaches. Such multi-agent systems can compartmentalize specialized task knowledge, organize themselves to avoid processing bottlenecks, and can be built flexibly to deal with dynamic changes in the agent and information-source landscape. Reported successful applications of multi-agent based approaches include organizational decision making, supply chain management, investment management, and military mission planning [32, 8, 20].

2. Many factors tender centralized systems to be inappropriate and sometimes infeasible for challenging problems in complex settings that are common in areas of Information Systems, Electronic Commerce, Enterprise Resource Planning, etc. For instance, data, knowledge, and decision-making in many applications are naturally distributed which makes a centralized approach inappropriate. In some other applications, centralized problem solving results in undesirable slow system responsiveness and a single point-of-failure. Prior to the advent of the World-Wide-Web and agent-based information management systems, researchers, primarily from the Artificial Intelligence (AI) field, have investigated many computationally difficult problems such as machine scheduling and resource allocation using distributed, agent-based approaches. These problems were originally formulated as centralized, classical decision or optimization problems. In many cases, however, centralized, non-agent based approaches are faced with serious difficulties (e.g., performance problems, inflexibility in dealing with changes in the environment, etc.), whereas agent-based formulation and problem solving approaches demonstrated their ability to deal with these difficulties more effectively [12]. In addition to some obvious advantages of agent-based approaches, such as natural representation of many inherently distributed problems, distributed computational characteristics of agent-based problem solving have proved to be highly desirable. Examples include system robustness, quickness in response to environmental changes, no single point-of-failure, etc.

3. There is a third category of agent-based applications that are important enough to be singled out. These are the applications that deal with *strategic* decisions tasks. Some examples are international trade negotiation, labor negotiation, and bargaining [14, 31, 16, 10, 24, 37]. These applications involve self-interested parties interacting with each other, trying to maximize their own payoff. For these tasks, often it is not possible to identify common goals among players, therefore tendering them fundamentally different from other decision problems which can be formulated as a single-agent problem (which, although, might be solved by a multi-agent, distributed approach).

   For non-strategic problems (such as the ones in the first two categories of agent applications), the overall system designers "own" the agents. The designers have complete control over the specifications and behaviors of agents including how these agents are constructed, what decision rules or knowledge that each agent has, what communication protocols are used, what coordination strategies are used by each agent to govern their interactions with other agents, etc. For strategic decision problems, this type of ownership and control no longer exists. Rather, agents serve their respective owners. They might have been built by different developers, representing potentially conflicting goals and agendas. For such systems, there is no place for overall system designers who make decisions regarding agents and agent coordination mechanisms. Agents will try to "outsmart" one another on behalf of their users. In addition, information received from other agents is no longer automatically trusted unless certain truth-telling mechanisms are enforced.

This type of agent-based applications is quite different from the first two types. Their practical significance is obvious in the new age of Information Economy (e.g., automated negotiation and auctions). Yet, very few such systems exist at the present time and even fewer have been deployed in real-world applications [1]. Most real-world systems oversimplify potentially important forms of strategic interactions (e.g., simple online auction bots).

We argue that the first and the third categories of software agents, i.e., "information management" and "strategic decision-making" agents are most relevant to procurement management given the particular characteristics and requirements of procurement tasks.

To take full advantage of e-commerce, collecting real-time information from a distributed collection of Web-based sources is essential. These sources often contain unstructured, multi-media, multi-model, and error-prone information, which is difficult to process. Post-processing and integrating collected data present another tedious and challenging task. Information management agents promise to alleviate user cognitive overload by automating information collection, integration, and continuous monitoring tasks. Limited success has been demonstrated in popular comparative shopping agents (e.g., mysimon.com and shopper.com). In our future work, we will develop and evaluate information agents for procurement. We will also explore issues relating to data mining and knowledge discovery on collection of past performance information. Research reported in this paper, however, ignores information management aspects of procurement. We simply assume the availability of the real-time flow of high-quality procurement information.

The application of the third category of agents to procurement constitutes our research focus. Typically the procurement process is complex and time-consuming. Current e-commerce technology provides real-time electronic collaboration infrastructures but rarely offers intelligent procurement decision-making or automated execution support which can be of tremendous value to procurement personnel in saving procurement time and efforts and also in ensuring the procurement quality.

In our research, we assume that the basic computer networking infrastructure exists to bring together participants involved in procurement-related transactions at a negligible cost. In addition, we assume that some (or all) participants will have access to automated procurement decision-making mechanisms, called procurement or trading agents, or simply, agents. Some of the simple agents are already widely used in B2C markets (e.g., proxy bidders from eBay). The more complex ones will be illustrated in Section 7. To avoid potential confusion concerning terminology, we reiterate our definition of agent in the context of this paper. Firstly, unlike in mainstream economics or business literature, we exclusively reserve the use of "agent" to refer to a computational entity. Secondly, with respect to the intelligent agent and multi-agent systems literature, we adopt an agent definition in a relatively weak sense [8] . In our context, an agent is simply any automated strategic procurement decision-making and execution system which satisfies the following set of conditions (with no reference to the level of decision-making sophistication or "intelligence"). (1) Agents operate in a networked environment. (2) Agents receive delegated procurement tasks from their human users. (3) Agents interact with other agents or human players directly or indirectly through well-defined online economic institutions. (4) Agents automate the overall process of procurement transactions. In Section 5, we discuss in detail the agent design and system development issues from a technical perspective.

Which kind of agent-based procurement decision-making support is appropriate to a large extent depends on the e-commerce environments and market rules or economic institutions in which procurement activities take place. In Section 3, we discuss a general research methodology based on experimental economics that can be leveraged to guide the development of and evaluate procurement agents. In Section 4, we focus on a specific type of online economic institutions based on auction that has great relevance to procurement practice.

3 Evaluating Procurement Technology using Experimental Economics

Procurement agents have great potential for improving the quality and efficiency of decision-making in the electronic marketplace and for reducing opportunity costs for parties participating in online procurement transactions. However, developing such agents requires addressing issues beyond the technical challenges of agent development (which will be the foci of Section 5). In order to guide the development of procurement agents and to assess their use and value, we apply experimental economics methodologies which enable in-depth analysis of variables of interest to online procurement in a controlled setting. In this section, we provide a brief survey of experimental economics methodologies and discuss their significant relevance to online procurement.

In general, experimental economics can provide economic principles and experimental designs for evaluating the IT infrastructure and agent support in trading and procurement contexts. It is useful to draw a distinction between laboratory experiments and field experiments. The methods for conducting laboratory experiments can be illustrated using an example from simple auction mechanisms.

Modern game-theoretic modeling of auctions originated in a classic paper by Vickrey [35]. He analyzed four types of auctions. The first type is the ascending-price auction in which the bidder with the highest bid wins the auctioned item at a price equal to his bid. This type of auction is called the English auction. The second type is called the Dutch auction in which the descending-price mechanism is used and the bidder with the first bid buys the auctioned item at a price equal to the bid. The third type is the first-price sealed-bid auction in which the bidder with the highest bid buys the auctioned item at a price equal to his bid. This type of bid is typically conducted in an off-line manner. The fourth auction type is the second-price sealed-bid Vickrey auction where the highest bid buys the auctioned item at a price equal to the second-highest bid.

Modern auction theory assumes that the bidders' values for the auctioned item are independently drawn from a known distribution, and that these values are private information. This can be implemented in the laboratory: the values can be independently drawn from a given distribution, and assigned to the subjects, who can be prevented from disclosing their values. To ensure that the bidders' economic behavior in the controlled experimental environment does not differ from that occur in real-world economic transactions, economic incentives are induced on them by paying the winner (high bidder) an amount of money proportional to the economic gain or profit which is equal to his value minus his bid. After running a series of auctions for each experimental condition, or a treatment, the subjects' bids and induced values for the abstract auctioned item can be used to analyze the effectiveness of the economic institutions under study and to examine the individual subjects' economic behavior under different institutions.

Smith in [28] explains the theory underlying the method of experimental economics and how microeconomics can be experimental science. For instance, by holding the environment constant in the laboratory but altering the institution, behavior may change such that the resulting market outcomes are markedly different. In fact, research in experimental economics has revealed that seemingly small modifications to trading rules can have prominent effects on competition. For example, early research in experimental economics has shown that double auction markets are nearly 100% efficient, in terms of achieving all available gains from trade, whereas posted offer markets are far less efficient. The New York Stock Exchange is an example of a double auction market, in which buyers and sellers submit bids and ask for stock ownership in firms. In contrast, retail markets for groceries or office supplies are posted offer markets, in which sellers post a price that buyers either accept or reject through their purchase decision. The inability of buyers to submit counteroffers to the sellers' prices leads to higher prices and lower efficiency, but saves each seller the transaction costs associated with setting price on a customer-by-customer basis.

Historically experimental economics has almost exclusively focused on the economics of human decision-making. New IT-enabled decision support systems such as software agents are an emerging phenomenon, with inextricable interactions with humans that will have economic implications for e-commerce and procurement management. Because the implementation of any new market design, especially one that includes autonomous purchasing mechanisms such as software agents, involves a number of key issues and a myriad of choices, we believe that the answers should come from an informed design. Markets on the Internet have a potential volume that was at one time unimaginable, and controlled laboratory experimentation can accelerate the process of designing attractive and effective intelligent procurement strategies in the electronic marketplace.

With an agent in the microeconomic system, a human expresses some behavior in endowing the agent with necessary parameters so that the agent can act on the person's behalf (the definition of an agent). The rules of the institution will still affect what actions the agent may take. However, one of the early findings in experimental economics research is that humans systematically under-reveal private information [29]. If people fundamentally under-reveal their private information in their own personal behavior; why would they not do they same when deploying agents? Thus, no matter how sophisticated the agent behavior, the layer of human-to-agent behavior plus the agent behavior may lead to different actions and market outcomes than to what actions and outcomes direct human behavior would lead. In a laboratory experiment, we can control agent behavior and observe human behavior towards agents. Moreover, we can compare these actions with the same individual's direct behavior in the same institution.

The other overlooked characteristic in agent-facilitated procurement is the mixing of human play with human-deployed-agent play. Direct human behavior and the concomitant tendency to under-reveal may change in a nontrivial way if mixed with human-to-agent play on the same side or other side of the market (buyer or seller). Our research attempts to examine the effects that human and agent behavior will have on efficiency and the distribution of surplus.

The following section outlines the specific procurement contexts in which we will explore (a) the layering of human behavior prior to the agent behavior, and (b) the human responses to agent play.

4 Web-enabled Procurement: Single-Sided and Doubled-Sided Auctions

An auction is a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants [17]. Auctions are of considerable interest to procurement management. The value of goods exchanged by auction is huge. Governments are the most prominent users of procurement auctions. For many government contracts, firms submit sealed bids; the contract is required by law to be awarded to the lowest qualified bidder. Auctions are also one of the most established trading mechanisms in the electronic marketplace. Electronic auction houses such as eBay are becoming part of everyday vocabulary. Impact of auctions on business-to-business procurement is being increasingly appreciated (e.g., freemarkets.com).

In addition to their obvious empirical significance, we have begun our work with auctions because auction mechanisms are relatively simple and well-studied in the literature. Agents for Web-enabled auctions are already being widely used. They are often called auctionbots. For instance, eBay provides automatic bidding systems. Bidders specify the maximum bidding amount that is kept private. The system then will automatically bid as the auction proceeds, bidding only enough to outbid other bidders. Although these bidding agents sound very straightforward, the interactions between agents and their human users are intriguing.

Four common types of single-sided auctions are either widely used or studied: the English auction, the Dutch auction, the first-price sealed-bid auction, and the Vickrey auction (see Section 3). Theoretically, all these four auctions for single indivisible units generate exactly the same expected profit for every (symmetric) seller and the same expected revenue for every buyer in a procurement setting under mild assumptions[18]. However, each auction has different properties that make it more desirable for different products. Different institutions are employed in traditional markets because of some of the inherent characteristics of the auctions, though many of these features are less relevant to the world of e-commerce. For example, the Dutch flower markets employ the Dutch clock auction because the highly perishable nature of the product mandates a speedy auction, the advantage of the Dutch clock auction. In the English auction, a bidder's optimal strategy does not depend on how competitors bid, and so it reduces the costs of information gathering and bid preparation. One common disadvantage of the Dutch clock and English auctions is that they require the actual presence of bidders, but with agents and e-commerce, this is no longer a prominent transaction cost. A second disadvantage of English auctions is that they are susceptible to bidding rings. (The members bid less aggressively and then split the profits among themselves in second auction.) While sealed-bid auctions do not require the presence of all the bidders, second-price auctions are also susceptible to fallacious behavior on the part of the person conducting the auction. After opening the sealed-bids, the auctioneer can insert his own false bid to manipulate the final price.

Part of our research focuses mainly on single-sided auctions in which a unique item is to be bought or sold. Specific issues and hypotheses under study include:

- Do people hide their true reservation prices from their bidding agents?
- Are agent-facilitated auctions more efficient both in terms of time and market efficiency than the traditional auctions which require direct human involvement?

In this study, we focus on existing, simple auctionbots. Our intention is to gain solid and thorough understanding of interactions among humans and simple auction agents that work on their behalf. We have examined these human-agent interaction issues in a series of auction scenarios. In Section 6, we present the experimental design, experimental findings, and analysis on agent-facilitated English auctions.

In addition to conducting evaluative research on simple auction agents for practical auction markets, we have also developed sophisticated agents that can coordinate biddings across a range of auction markets, including both single-sided and double-sided auctions. As opposed to single-sided auctions, in a double-sided auction, several buyers and several sellers submit bids simultaneously. Evaluation of such sophisticated agents are conducted through competition between them. Section 7 motivates the development of such agents and provides detailed technical discussion concerning related intelligent trading strategies.

5 An Agent-enabled Procurement Testbed

After surveying currently implemented e-commerce and procurement applications and testing environments, including both commercial systems and research prototypes, we conclude that existing systems do not provide adequate control and flexibility needed to carry out the research laid out in the previous sections. Thus, as part of our EARP effort, we have developed a flexible e-commerce testbed leveraging existing e-commerce research and our systems experience on knowledge-based systems, Web-enabled applications, and software agents.

Our testbed is designed to interact with both human subjects and trading agents working on behalf of their users. The testbed has been structured such that it can be used in two modes. In the first mode, called the "synchronous" mode, the testbed is used in a controlled laboratory setting where human subjects are be physically present and participate in well-structured experiments through assigned computer terminals. In the second mode, called the "asynchronous" mode, human subjects can use Web browsers from their own computers to participate in the experiments in a distributed manner and can interact with the testbed as they normally would with commercial e-commerce sites.

To ensure maximum system compatibility and extensibility, we built this flexible online procurement testbed in the Java programming language. As illustrated in Figure 1, our testbed has two major components: (a) an auction server that implements various economic institutions under study, and (b) a set of automated trading mechanisms implemented as procurement agents, accessible to the user through Java Applets.

The auction server manages user registration and maintains the list of objects that are being traded. All transactions and message exchanges are logged on the server. The standard agent messaging format based on KQML [5] is used in our system and protocol design. The main purpose of the server is to provide services regarding the "rules of games." Configurable by the experimenters, the server controls the institutional rules such as when and how to clear the market, what information is made available to the traders, etc. It also provides a graphical interface for the experimenters to monitor the progress of ongoing auctions.

The second major component of the testbed consists of a set of procurement or trading agents from which experimenters or human subjects can choose. These agents co-exist

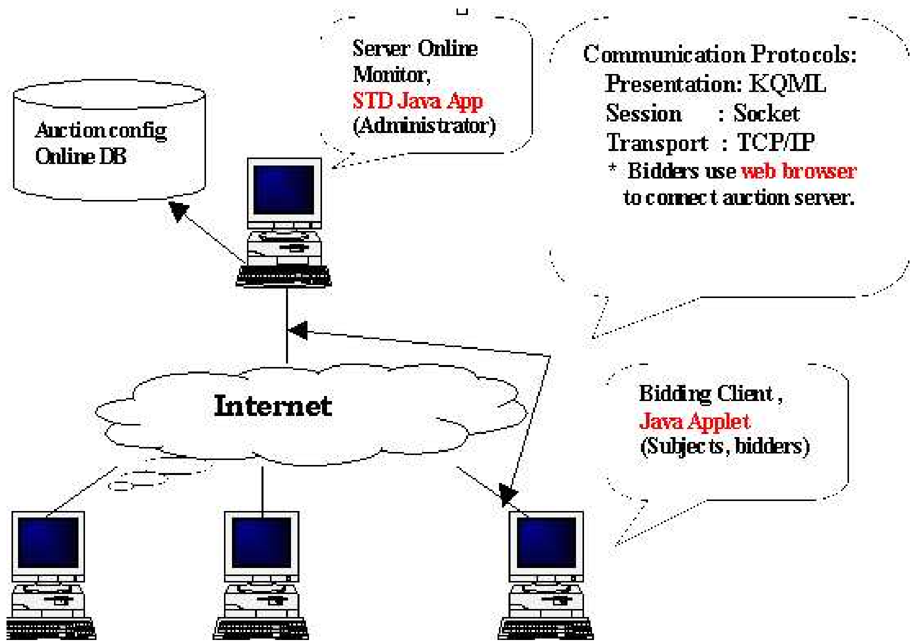

Figure 1: System Architecture

and interact with human subjects, enabling us to study the impact of mixing people and software agents in economic transactions. Code of such agents resides on the auction server. When a bidder accesses the server to join in an auction through a Web browser, an agent, implemented as a Java applet, is first downloaded to the bidder's client machine and then starts execution via establishing communications with the server. After that, a bidder will be participating in the auctions through this agent. Depending on the type of the agent (selected either by the bidder or forced by the experimenter), the bidder can specify the specific auction behavior that the agent will manifest and then let the agent automate the auction process under strict guidance. Or the bidder can initiate intelligent trading strategies where the agent has higher degree of autonomy with less human involvement.

Our testbed is structured such that two types of agents are supported.

1. *Simple execution agents.*

   This type of agent does not have an explicit model of the underlying procurement task. Nor does it provide active decision support. The sole functionality of a simple execution agent is to automatically execute its human user's trading strategies that are fully specified by the user at the beginning of trading through a limited number of parameters. Popular proxy bidding agents (e.g., the eBay proxy bidder) fall into this category. These agents are initialized with their user's reservation price. Then without user intervention, they will monitor the progression of an auction and place a bid when the current highest bid is below the pre-specified reservation price and is held by another agent. Our evaluative study reported in Section 6 focuses on this type of simple agents.

2. *Behavioral agents.*

   This type of agent possesses explicit knowledge of the trading institutions and provides active decision support. Computational models of auctions and negotiations from the Economics and AI literature can serve as the basis for constructing such agents. To be computationally feasible, these models typically ignore certain aspects of strategic interaction. An example of such agents is presented in Section 7.

We use a sequence of screen-dumps to illustrate an example operation of the developed testbed as in the "synchronous" or experimentation mode. Figure 2 shows the interface that an experimenter uses to specify an auction session which is defined as a set of closely-related auctions that a set of bidders will participate in sequentially in one batch. Through the GUI, the experimenter can specify the auction items, induced values for each bidder, and whether or the way through which the bidders will be allowed to access procurement agents. The experimenter can also specify whether the very fact that a bidder is using an automated mechanism can be made public to other bidders. To facilitate experimental design, the experimenter can set up all these parameters in an Excel sheet and import them to the experimental setup utility without manual input through the GUI.

After a bidder logs on to the auction server, she first provides information regarding her past auction experience and other demographic information through the interface shown in Figure 3.

After user registration, a procurement agent is downloaded from the auction server and the bidder interacts with the agent through an interface shown in Figure 4. This interface provides information concerning the auction item and detailed auction rules (e.g., the minimum bid increment between bids). It also shows bid activity and history which are empty initially. Furthermore, it shows time remaining in the auction. The most important part of the interface is the choices that the bidder can choose from regarding the operation mode of auction (e.g., between manual bid or proxy agent bid).

Figure 5 shows another snapshot of the same interface. Note that this interface corresponds to an ongoing auction instance, populated with bid activity and history

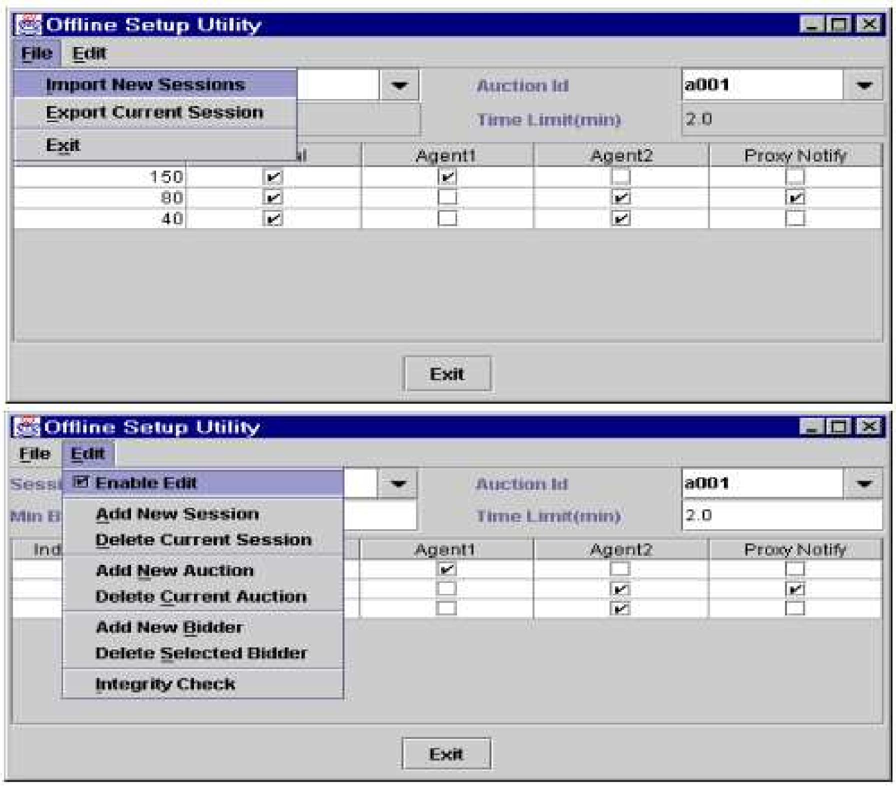

Figure 2: Experimental Setup

Figure 3: Bidder Registration

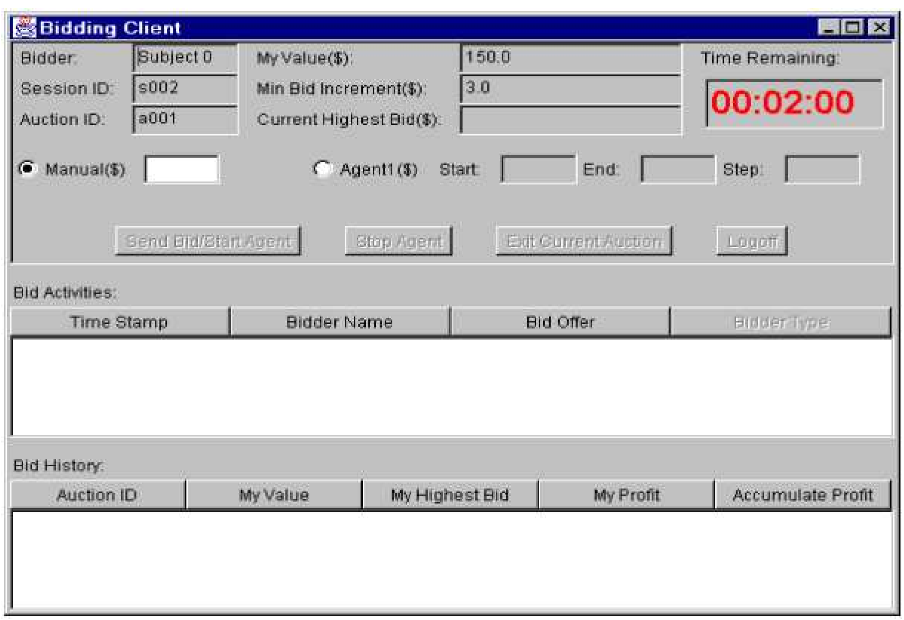

Figure 4: Bidder Screen (1): Before The Auction Starts

data.

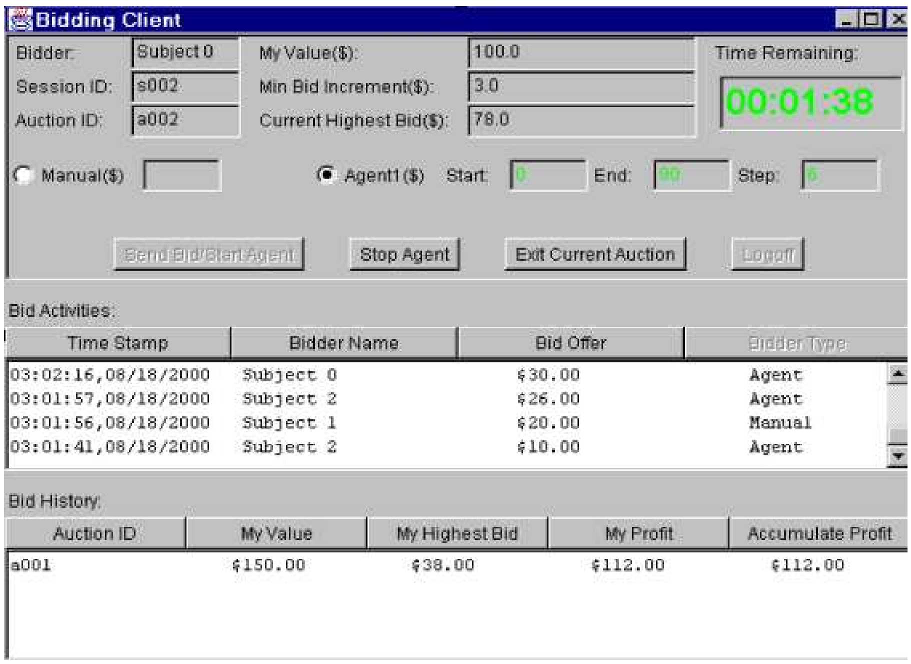

Figure 5: Bidder Screen (2): Ongoing Auction

To facilitate the experimenter's control over and monitoring of ongoing auctions, a separate utility has been implemented, as illustrated in Figures 6 and 7, corresponding to the status of an auction before and after it starts, respectively.

We summarize the key differences between our testbed and existing systems. (1) Our system provides both "synchronous" and "asynchronous" operations, enabling us and other researchers to conduct both laboratory experiments as well as online field studies. Almost all existing systems deal only with "asynchronous" operations and do not provide adequate support for laboratory experiments. (2) Our testbed comes with a set of automated trading strategies encapsulated as agents. Most other systems only provide the e-commerce server services. (3) Our testbed enables agent and human participation simultaneously. This capability makes it possible for us to systematically study various forms of human-agent interactions. With the exception of AuctionBot, other systems typically focus on either human participation or agent participation but not both.

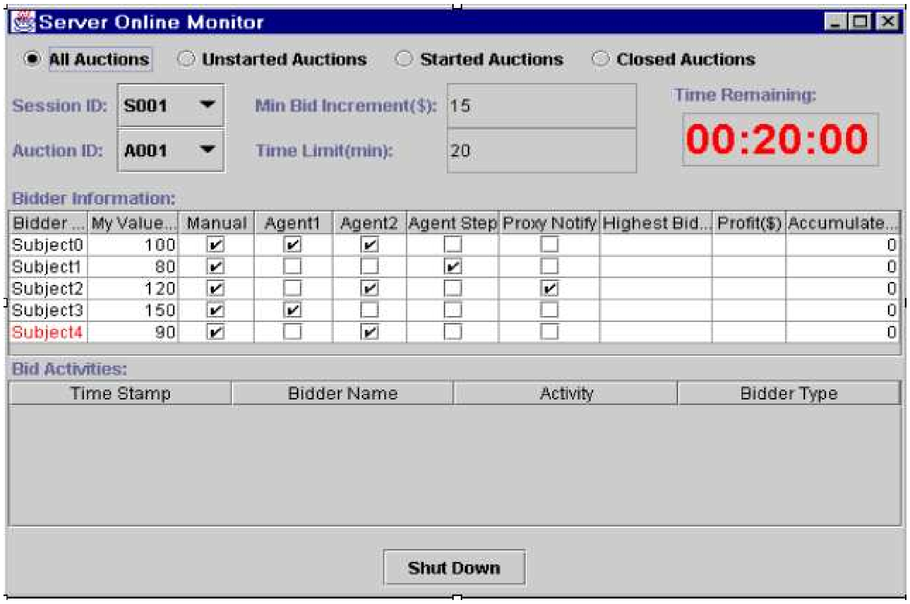

Figure 6: Experimenter Screen (1): Before The Auction Starts

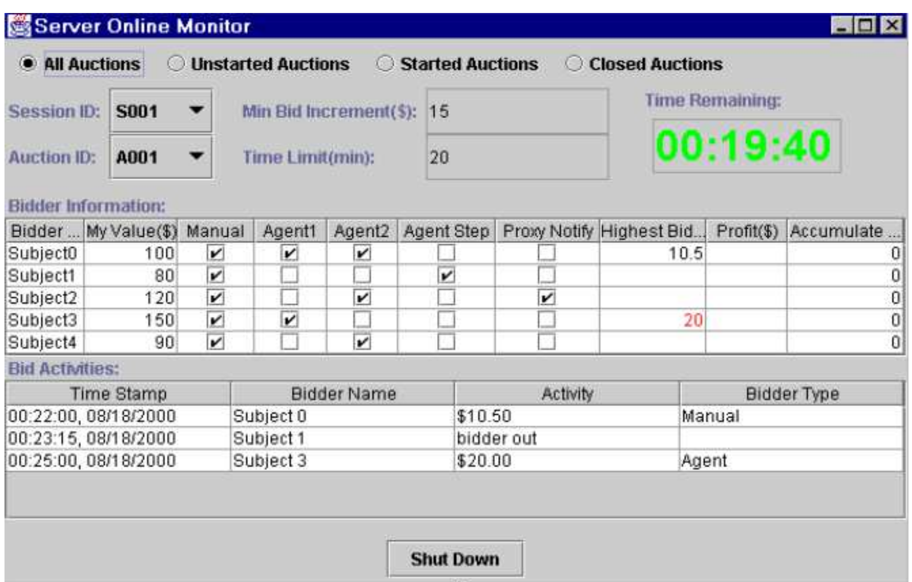

Figure 7: Experimenter Screen (2): Ongoing Auction

6 An Experimental Study: Web-based English Auctions

Using the online procurement testbed described in the previous section, we initiated research to examine how the layering of human behavior prior to the agent behavior affects market performance. As a first step we conducted an experimental study to explore human behavior associated with a prevalent, but primitive agent in Internet auctions. This simple experiment illustrates how simple agent behavior may affect auction markets.

Commercial auction sites such as eBay and Yahoo!Auctions offer their bidders the option of submitting proxy bids on their behalf. The proxy bidding system is fairly simple in an English auction. The bidder specifies a maximum bid and the website will automatically outbid the highest current bid by the minimum increment, up to the maximum amount specified by the bidder. Because the maximum bid is private, the dominant strategy in an English auction is to submit to your value as the maximum bid. To explore the behavior of markets with a proxy bidding agent, we recruited 20 University of Arizona undergraduates to participate in a one hour experiment using Java applet auction software that we developed. The student subjects each participated in 24 auctions, two of which were running concurrently. Each market consisted of four bidders.

As is standard laboratory auction experiments, the subjects were informed that if they purchased a unit for less than their value, they would earn the difference as profit, which would be converted into cash at a rate of 8 experimental dollars to one US$.[1] One bidder in each market was first instructed to submit a single proxy bidding agent to the market. Once submitted, the remaining three bidders were free to bid in real time for the next two minutes. By forcing the proxy bidder to bid first and only once, we are introducing an explicit transaction cost for real time monitoring and bidding in the electronic auction. (Proxy bidding was created to reduce such transaction costs.) The minimum increment in the auction was 1. The highest value for each auction was randomly drawn from a $U[90,110]$ distribution, the second highest from a $U[80, 94]$ distribution, the third highest from a $U[69, 83]$ distribution, and the lowest from a $U[60, 74]$ distribution. The values for each bidder were held constant across the five market sessions, and the items were sequentially auctioned, two at a time, in each market. Since we are interested in

[1] In addition to their earnings, subjects received $5 for showing up on time. The average earnings without the show up fee were $12.75. Hence, $17.25 was more than sufficient to cover the opportunity cost of one hour's time for an undergraduate student.

the layering of human behavior prior to the agent behavior, the subject with the proxy bidding agent was given the highest value drawn for 12 of the 24 auctions. The remaining bidders each received the highest value twice and were divided among the other three relative positions.

The theoretical prediction is that the highest value bidder will win the auction paying a price equal to the second highest value. Early experimental work in the late 1970's and early 1980's focused on behavior in auctions. We define efficiency as the value of the winning bidder divided by the highest bidder. Hence, an auction is 100% efficient if the highest value bidder wins the auction, i.e., the gains from trade are maximized. Coppinger, Smith, and Titus found that English auctions were nearly always 100% efficient, but with prices slightly greater than the second highest value when the minimum increment is not very small [3].

We summarize our experimental findings in Figure 1 and Table 1. Figure 1 illustrates the average winning bid for the five sessions and the optimal price (the second highest value). We observe that the presence of a proxy bidder consistently leads to prices on average less than the second highest value. One reason for the lower prices is that a subject specifies a maximum bid for his agent that is less than his true value, when the subject has the highest value. When this maximum bid is less than second highest value, the bidder with second highest value wins the auction. In other auctions, even when the proxy bidder does not have the highest value, the real time bidders tire from being outbid by 1, the minimum increment, by the proxy bidder, and then they give up bidding. They wait for the final seconds of the auction to tick off, and then in the final seconds, the bidders submit bids well below their values so that the highest value bidder again may not win. In essence, the presence of the proxy bidding agent turns the English auction into a last second first price auction.[2] The implication of this behavior is that the auctions can be inefficient. Table 2 reports that efficiency levels of the sessions are 4.2 to 8 percentage points below the 100% level. This is due in large part to the fact that the proxy bidder, who has the highest value in 12 of the 24 auctions, only wins 5.6 times on average (out of 12). As a result, the seller (experimenter) receives a price less than the price at the 100% efficient level (78.6 vs. 85.4).

We must caution that our results are preliminary due to the small size of our sample. We are currently working on running baseline sessions without proxy bidders to verify that the English auction is 100% efficient. We also need to test the robustness of these results

[2]Our English auction closed after 2 minutes and 30 seconds. The subject with the proxy bidding agent was given the first 30 seconds to submit his agent. As soon as the proxy bidding agent was deployed, the real time bidders could bid with the remaining time on the clock.

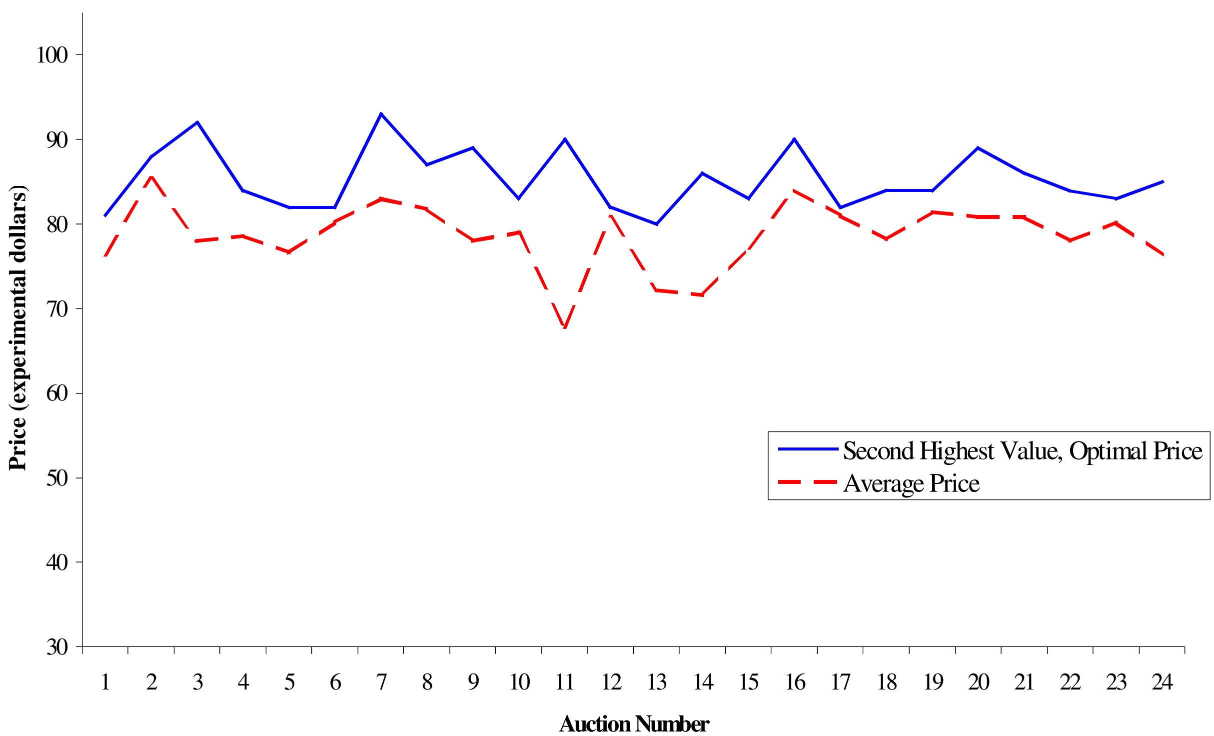

Figure 8: Average Price per Auction (5 sessions)

with more than 4 bidders. Nevertheless, this experiment illustrates that introducing simple agents into an auction may affect market performance.

Session #	1	2	3	4	5	All
Average Efficiency	92.0%	95.1%	95.1%	95.8%	93.8%	94.4%
# of Proxy bidder Wins	2/12	7/12	9/12	5/12	5/12	5.6
Average Price	73.3	82.0	72.7	84.7	80.5	78.6
100% Efficient Price	85.4	85.4	85.4	85.4	85.4	85.4

Table 1: Summary Statistics

7 Developing Intelligent Trading Strategies for Multi-Item Auctions

Single-item auction mechanisms and related human-agent interaction research issues described in previous sections are important for online procurement research. Yet, the real-world procurement practice typically involves more complex mechanisms that need to be systematically studied. In this section, we present an important extension to the single-item auction research that involves purchasing of several related items from different markets simultaneously. When multiple items have to be purchased as a *bundle*, the value of one item may depend on what other items he can obtain and as a result, auction activities have to be coordinated in a close manner.

This section focuses on an intelligent trading strategy that we have developed in making real-time bidding decisions across a spectrum of disparate markets. In this study, evaluation of our approach is performed through an agent tournament. Because of the complexity of the procurement task involved, as opposed to the study reported above, we do not address human-agent interaction issues. Rather, we emphasize computational aspects of auction strategies and study agent-agent interactions.

We first provide background information on the procurement task which motivated our study. The first international Trading Agent Competition (TAC-2000) challenged its entrants to design an automated trading agent that was capable of bidding in simultaneous on-line auctions for complementary and substitutable goods [6]. A TAC agent is a simulated travel agent whose task is to organize itineraries for a group of clients who wish to travel from TACTown to Boston and back again during a five-day period in July. Travel goods, such as airline tickets and hotel reservations, are complementary, and tickets to entertainment events, such as the Boston Red Sox and the Boston Symphony Orchestra, are substitutable. The trading agent's objective is to win items that best satisfy its clients' preferences as inexpensively as possible. The goals of the tournament included providing a benchmark problem in the complex domain of e-marketplaces, and motivating researchers to apply unique approaches to a common task.

The general bidding strategy employed by the travel agent that we have developed, called UATrader, can be characterized as a "myopic" trading strategy with iterative adjustments based on neighborhood search. After analyzing several game instances and observing the behavior of some participating trading agents, we hypothesized that in most cases, the decisions regarding what days each customer should stay in Boston play a deciding role in agent performance. Based on this hypothesis, UATrader actively participates in the flight and hotel auctions and seeks to coordinate its bidding activities in these auctions, while its involvement in the entertainment ticket auctions is secondary and largely opportunistic.

UATrader bases most of its bidding decisions on the *anchor solution*—a hypothetical assignment of travel dates for each customer. The anchor solution is initialized to reflect each customer's preferred arrival and departure dates. Whenever a price change is observed, UATrader evaluates small variations (neighbors) of the current anchor solution on a customer-by-customer basis. This evaluation is based on the sum of the potential changes across all three types of auctions given the current price quotes. If the overall impact of switching from the current anchor solution to a variation is positive, this variation is made the new anchor solution for the corresponding customer.

UATrader bids in two phases. The first phase lasts for the first 13 minutes of the game. It switches to the second phase for the remaining 2 minutes.

Flight Auctions. In the first bidding phase, UATrader submits fixed low bids (e.g., $165) for all available flights to take advantage of possible low airfares. In the second phase, UATrader tries to assure the booking of the flights that match the anchor solution for each customer by submitting high bids.

Hotel Auctions. UATrader's behavior in hotel room auctions is closely coordinated with its bidding strategy for flights. In the first phase, UATrader submits dummy bids at fixed time intervals to keep the auctions from closing. In the second phase, UATrader relies on anchor solutions to examine whether the Grand Hotel (BGH) or Le Fleabag Inn (LFI) rooms should be targeted. Initially, UATrader assumes that all customers should stay in BGH. Then, as the auctions proceed, for each customer, UATrader estimates the utility change of switching from BGH to LFI (or switching from LFI to BGH if the current target is LFI) based on the current prices. If for certain nights, prices of both hotels exceed pre-specified thresholds, UATrader automatically modifies the corresponding anchor solutions and adjusts its bids in the flight auctions.

Entertainment Ticket Auctions. In the first bidding phase, UATrader is not active in any of the entertainment ticket auctions. In the second phase, after the anchor solution is booked, UATrader makes a one-time decision based on the current price quotes as to selling and buying tickets with the objective of maximizing the total customer utilities.

UATrader participated in the qualifying rounds of TAC-2000 and TAC-2001 and won a seat in the final rounds of both competitions. Its performance is generally satisfactory. We are currently in the process of analyzing the interactions between UATrader and other trading agents developed by other research teams.

8 Summary and Future Research

In this paper, we bring together advanced information management and decision making technologies such as software agent and rigorous experimental economics design and methodologies to make contributions to procurement and e-commerce research.

The potential impacts of this inherently multi-disciplinary research are twofold:

- **Benefiting e-commerce, procurement and experimental economics research communities.** The research impacts procurement related IT research communities by incorporating the science of experimental economics as both a structured evaluation framework and an economic content provider.

- **Offering managerial insights and practical guidance for practitioners.** Our research generate technological and managerial insights on identifying right mix of "IT" and economic institutions for effective and efficient procurement practice.

Several follow-up research projects and extensions to the reported work are being actively pursued.

Firstly, we are currently experimenting with other auction types besides the English auction. The online procurement testbed provides a flexible technological base for such experimentation. The set of research issues to be addressed is similar to that for the English auction: the economic impact of introducing simple automation into auctions, and human-agent interactions.

Secondly, we are using the procurement testbed to conduct field experiments. It is more difficult to interpret data from field experiments or non-experimental markets because bidders' auctioned item values are not observable. However, the impact of IT and the effects of changing the market institution can be tested with field data given controls on the distribution of auctioned item values (following, e.g., [13]).

Thirdly, we are extending the approach implemented in UATrader such that it can be used in a generic procurement setting rather than just for vacation packages. Research has been planned to validate this agent approach that demonstrates sophisticated auction behaviors and coordinative actions across markets via both simulation and controlled experimentation.

Acknowledgments

Additional funding for the research results reported in this paper is provided by the Faculty Small Grants Program sponsored by The University of Arizona Foundation, and by the International Foundation for Research in Experimental Economics at the University of Arizona. Dr. Bart Wilson is a key collaborator and has contributed his expertise in experimental economics to our project. He is a co-author for the journal submissions based on reported research. We would like to thank Drs. James Cox and Hsinchun Chen for insightful discussions concerning experimental economics and e-commerce. We also wish to recognize the following University of Arizona students who have helped in the auction testbed system implementation: Jiang Zhu, Pengfei Wang, and Pei He.

References

[1] K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, 1(1–9), 1999.

[2] A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, 1996.

[3] V. Coppinger, V. Smith, and J. Titus. Incentives and behavior in english, dutch, and sealed-bid auctions. *Economic Inquiry*, 18(1):1–22, 1980.

[4] O. Etzioni and D. Weld. A softbot-based interface to the internet. *Communications of the ACM*, 37(7), July 1994.

[5] T. Finin, R. Fritzson, and D. McKay. A language and protocol to support intelligent agent interoperability. In *Proceedings of the CE and CALS Washington 92 Conference*, June 1992.

[6] A. Greenwald and P. Stone. Autonomous bidding agents in the trading agent competition. *IEEE Computing*, 2001.

[7] N. Jennings, P. Faratin, T. Norman, P. O'Brien, B. Odgers, and J. Alty. Implementing a business process management system using adept: A real-world case study. *International Journal of Applied Artificial Intelligence*, 2000.

[8] N. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi Agent Systems*, 1(1), July 1998.

[9] A. Joshi and M. P. Singh. Multiagent systems on the net. *Communications of the ACM*, 42(3), 1999.

[10] S. Kraus and J. Wilkenfeld. A strategic negotiation model with applications to an international crisis. *IEEE Transaction on Systerms, Man, and Cybernetics*, 23(1):313–323, 1993.

[11] N. Krothapalli and A. Deshmukh. Design of negotiation protocols for multi-agent manufacturing systems. *International Journal of Production Research*, 37(7):1601–1624, 1999.

[12] J. Liu and K. Sycara. Coordination of multiple agents for production management. *Annals of Operations Research*, 75:235–289, 1997.

[13] D. Lucking-Reiley. Using field experiments to test equivalence between auction formats: Magic on the internet. *American Economic Review*, 89(5):1063–1080, 1999.

[14] M. Ma. Agents in e-commerce. *Communications of ACM*, 42(3):78–80, March 1999.

[15] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7), July 1994.

[16] P. Maes, R. Guttman, and A. Moukas. Agents that buy and sell. *Communications of ACM*, 42(3), March 1999.

[17] R.P. McAfee and J. McMillan. Actions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.

[18] P. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, Summer 1989.

[19] T. Mitchell, R. Caruana, D. Freitag, J. McDermott, and D. Zabowski. Experience with a learning personal assistant. *Communications of the ACM*, 37(7), July 1994.

[20] M. Nissen and A. Mehra. Some intelligent software supply chain agents. In *Proceedings of Agent 99*, May 1999.

[21] Y. Peng, T. Finin, Y. Labrou, R. Cost, B. Chu, J. Long, W. Tolone, and A. Boughannam. Agent-based approach for manufacturing integration: The ciimplex experience. *Applied Artificial Intelligence*, 13:39–63, 1999.

[22] R. Quick and E. White. E-business: E-retailers now facing holiday sales slowdown. In *Wall Street Journal*, Dec 11 2000.

[23] D. Riecken. Intelligent agents. *Communications of the ACM*, 37(7):18–21, July 1994.

[24] J. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, Cambridge, Mass., 1994.

[25] T. Sandholm. emediator: A next generation electronic commerce server. In *International Conference on Autonomous Agents (AGENTS)*, Barcelona, Spain, June 2000.

[26] M. Shaw and R. Sikora. A multi-agent framework for the coordination and integration of information systems. *Management Science*, 44(11), 1999.

[27] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.

[28] V. Smith. Microeconomic systems as an experimental science. *American Economic Review*, 72:923–955, December 1982.

[29] V. Smith. *Papers in Experimental Economics*. Cambridge University Press, 1991.

[30] J. Swaminathan, S. Smith, and N. Sadeh. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences*, 29(3), summer 1998.

[31] K. Sycara. Problem restructuring in negotiation. *Management Science*, 37(10):1248–1268, 1991.

[32] K. Sycara and D. Zeng. Coordination of multiple intelligent software agents. *International Journal of Cooperative Information Systems*, 5(2 & 3):181–211, 1996.

[33] M. Totty. E-commerce (a special report): Overview—the next phase. In *Wall Street Journal*, May 21 2001.

[34] D. Veeramani, P. Joshi, and V. Sharma. Critical research issues in agent-based manufacturing supply webs. In *Proceedings of Agent-based Decision-Support for Managing the Internet-Enabled Supply-Chain Workshop, Autonomous Agents 1999*, pages 109–118, 1999.

[35] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *J. Finance*, 16:8–37, 1961.

[36] M. Wooldridge and N. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.

[37] D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141, 1998.